A Non-Intrusive and Context-Based Vulnerability Scoring Framework for Cloud Services

Hao Zhuang, Florian Pydde
EPFL

arXiv:1611.07383v2 [cs.CR] 7 Dec 2016

ABSTRACT

Understanding the severity of vulnerabilities within cloud services is particularly important for today's service administrators. Although many systems, *e.g.*, CVSS, have been built to evaluate and score the severity of vulnerabilities for administrators, the scoring schemes employed by these systems fail to take into account the contextual information of specific services having these vulnerabilities, such as what roles they play in a particular service. Such a deficiency makes resulting scores unhelpful. This paper presents a practical framework, NCVS, that offers automatic and contextual scoring mechanism to evaluate the severity of vulnerabilities for a particular service. Specifically, for a given service *S*, NCVS first automatically collects *S*'s contextual information including topology, configurations, vulnerabilities and their dependencies. Then, NCVS uses the collected information to build a contextual dependency graph, named CDG, to model *S*'s context. Finally, NCVS scores and ranks all the vulnerabilities in *S* by analyzing *S*'s context, such as what roles the vulnerabilities play in *S*, and how critical they affect the functionality of *S*. NCVS is novel and useful, because 1) context-based vulnerability scoring results are highly relevant and meaningful for administrators to understand each vulnerability's importance specific to the target service; and 2) the workflow of NCVS does not need instrumentation or modifications to any source code. Our experimental results demonstrate that NCVS can obtain more relevant vulnerability scoring results than comparable system, such as CVSS.

1. INTRODUCTION

Today's cloud-scale systems become increasingly complex – they not only employ multi-layered network/software stacks, but also deploy various distributed service components offered by other providers. These structurally complex systems, nevertheless, may inadvertently introduce much more vulnerabilities than traditional computer systems.[1] Therefore, it is important for administrators to score vulnerabilities based on their severity, thus enabling administrators to deal with the critical ones accordingly.

Driven by the above motivation, many public vulnerability databases (*e.g.*, CVE [1] and OSVDB [5]) are maintained, which contain classification of vulnerabilities, description of the nature of their severities, as well as the exploitability based on the feedback of security experts. Furthermore, several vulnerability scoring systems are also developed for the disclosure and severity ranking of vulnerabilities. One of the most representative vulnerability scoring systems is Common Vulnerability Scoring System (CVSS) [2], which is designed based on expert knowledge, multi-dimension scoring schemes and customer feedback.

However, current vulnerability scoring efforts typically quantify the severity of vulnerabilities in a general (or global) way – *i.e.*, how severe vulnerabilities are – rather than considering the context of specific services holding these vulnerabilities. Applying such scores to evaluate the risk of a particular service, therefore, may not always be meaningful or instructive. For example, a denial of service vulnerability against MySQL might have a high severity score in CVSS. But for a particular service context, its application components running MySQL may only be used for exception event logging, which is rarely invoked. Thus, the impact of this vulnerability on the whole service context would not be as significant as the ranking score suggested in CVSS. On the contrary, if a service context has to rely on such legacy systems for its core functionality, their vulnerabilities should be ranked high and handled with high priority.

Based on this insight, for a particular service administrator, it is desirable to have a tool, which not only considers the properties of vulnerabilities, but more importantly takes into account the target service's context. Such a tool can offer a truly meaningful severity score for each vulnerability to serve as a guideline regarding which vulnerabilities are the most urgent and critical ones for that *particular service*, thus enabling administrators to allocate appropriate patches or fixing code to deal with them accordingly. Although many efforts (*e.g.*, attack graph approaches [13,15,24,33] and profiling techniques [34, 38]), these efforts are either *ad hoc* or hard to be extended to tackle our problem (see §2.2 for more details).

This paper presents a systematic work, NCVS, which is a practical contextual vulnerability scoring framework for cloud services. To evaluate the severity of vulnerabilities in terms of a particular service context, NCVS first automatically collects contextual information about service components and their dependencies in a comprehensive way. Using this acquired data, NCVS then builds a *contextual dependency graph*, named CDG, to model the service's context. Finally, NCVS analyzes all the vulnerabilities' relative importance in this context and returns a contextual vulnerability scoring report. To the best of our knowledge, NCVS is the first effort capable of offering automatic (*i.e.*, non-intrusive) and context-aware vulnerability scoring for cloud

[1] A vulnerability is a weakness (defect) of the software design or implementation rather than necessarily a bug; in other words, a system may have a vulnerability due to a defective design, even if its implementation is perfect.

services.

Building a non-intrusive system for contextual vulnerability scoring needs to address several challenges. First, the information of components in the target cloud service and their associated dependencies should be comprehensively acquired, since vulnerabilities may exist in different types of components. In addition, infrastructures underlying today's services tend to be complex; thus, asking an administrator to manually collect such a large dataset is an infeasible task. Existing techniques *either* heavily rely on interruptive instrumentations (*e.g.*, MagPie [9] and Project 5 [7]), *or* have been developed for limited types of dependencies (*e.g.*, Sherlock [8] and NSDMiner [14, 20]). Thus, how to *comprehensively* acquire the detailed contextual dependencies (including network, hardware and software-level dependencies) *without instrumentations* presents a challenge. We develop a non-intrusive approach capable of collecting three types of contextual information by mining log information and network traffic (detailed in §3.2).

The second challenge includes 1) how to model the context with the collected information and 2) how to compute the importance of each vulnerability by taking into account this context. We propose a new representation, named contextual dependency graph (or CDG), to model the target service's context. In general, CDG is a two-layered directed acyclic graph (DAG) representation, and is capable of capturing more details than previous dependency graph models. With CDG in hand, we propose a pluggable scoring module that allows administrators to apply diverse (existing or customized) ranking algorithms based on their requirements. As an example, we also propose a new algorithm specific to our CDG model by adapting PageRank algorithm [23]. Our experimental results indicate the new algorithm can obtain more reasonable vulnerabilities' scores than CVSS (detailed in §3.3).

In summary, this paper mainly makes four contributions. First, we present the first practical (*i.e.*, general, non-intrusive, pluggable and effective) contextual vulnerability scoring framework for cloud services. Second, we develop a set of collectors capable of automatically acquiring a target cloud service's contextual information, including network, service and hardware dependencies. Third, we construct a new graph model, named CDG, specific to our collected information and vulnerability scoring purpose. Finally, we demonstrate the practicality of NCVS based on a lab-scale case study and performance measurements.

2. MOTIVATION AND RELATED WORK

This section first motivates our work (§2.1), and then discusses why existing work does not help (§2.2).

2.1 Motivating Example

Figure 1 presents a simple but illustrative example for our motivation. In this example, we consider a real-world lab-scale Hadoop cluster with eight servers: one name node (S1), one backup name node (S5), and six data nodes (S2-S4, and S6-S8). The name nodes and data nodes run job trackers and task trackers, respectively. S1-S4 and S5-S8 belong to two individual racks, respectively. Each rack has one edge switch [29, 37].

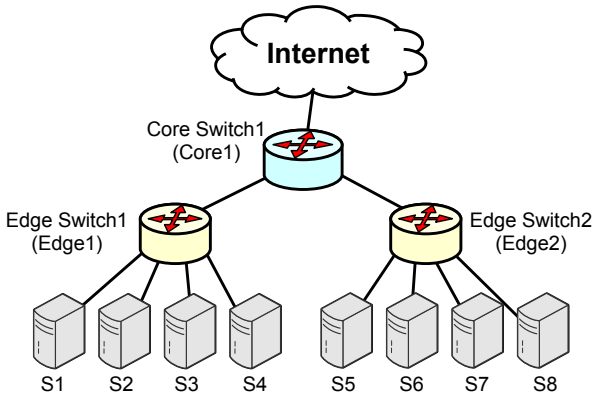

Figure 1: Motivating example.

By retrieving one of the most representative vulnerability scoring systems, Common Vulnerability Scoring System (CVSS) [2], the administrator discovers two vulnerabilities in this cluster: 1) Hadoop name node vulnerability, CVE-2015-7430, with severity score 8.4; and 2) core switch vulnerability, CVE-2016-1392, with severity score 7.4. In the example of Figure 1, the former exists in Hadoop components running in S1 and S5, while the later exists in Core1 switch software. If this administrator prioritizes her fixing or patching vulnerabilities based upon the scores provided by CVSS, she should fix Hadoop vulnerability first. Nevertheless, if we deeply analyze the given Hadoop cluster, we may find that the states of name nodes in this Hadoop cluster have been replicated across two servers (*i.e.*, S1 and S5). Even though the vulnerability CVE-2015-7430 is triggered on one name node (*e.g.*, S1), the service can use S5 – the backup name node – to manage all data nodes. On the contrary, if the vulnerability CVE-2016-1392 is triggered, any network traffic from external network cannot reach any servers in this cluster.

Based on the above example, we note that existing vulnerability scoring systems fail to guide administrators to understand the severity of vulnerabilities in their service deployments. In other words, because existing vulnerability scoring systems do not take into account contextual information of target vulnerabilities, severity scores offered by these systems are not truly helpful in practice. Thus, we raise a question that whether it is possible to build a practical contextual vulnerability scoring system?

This paper develops such a system, named NCVS. In Figure 1 example, if the administrator uses NCVS to score the two revealed vulnerabilities, she can obtain a different result: CVE-2016-1392 has a higher severity score than CVE-2015-7430 in terms of service context as shown in Figure 1.

2.2 Existing Efforts Discussions

We now describe and compare existing efforts, and discuss why they do not work for our purpose.

Vulnerability scoring systems. Many current vulnerability scoring systems (*e.g.*, CVSS [2], OSVDB [5] and OVAL [4]) have been developed to score vulnerabilities based on their severity. However, these systems only score vulnerabilities generally, *i.e.*, how severe they are by themselves without considering any specific environment. Such a score is not meaningful or instructive to evaluate vulnerabilities in specific service context (like the example in §2.1).

Although several contextual vulnerability scoring efforts, *e.g.*, improved CVSS [3, 11] and VRank [15], have been proposed, they are quite *ad hoc* and impractical. First, these efforts do not have any automatic environmental information collection capability. This means all the contextual information has to be input manually, making their scoring process impractical to any complex services. Second, these approaches only support coarse-grained scoring schemes, because they lack expressive models. Finally, the environmental factors considered by these efforts are too simple to represent the contexts of real-world services.

Attack graph based risk evaluation efforts. Attack graph based risk reasoning approaches [13, 21, 22, 24] have been studied in the past ten years. An attack graph is an abstraction of the details of possible attacks against any computer system or service. Attack graph based efforts are appropriate to evaluate the potential risks of target services, rather than computing vulnerabilities' severity. In other words, attack graph techniques leverage existing vulnerability scores (*e.g.*, provided by CVSS) to compute the risks of services of interest, but it is not capable of updating vulnerabilities' scores based on different services' contexts.

3. NCVS DESIGN

This section first presents a high-level overview of NCVS framework in §3.1. Then, we detail two important modules in §3.2 and §3.3, respectively.

3.1 NCVS Framework Overview

As shown in Figure 2, NCVS has two important modules: 1) contextual information acquisition module and 2) vulnerability scoring module. In addition, NCVS relies on two databases – contextual information DB and vulnerability threats DB – that are used to provide necessary information for vulnerability scoring.

NCVS performs the following two steps to score vulnerabilities for a target cloud service *S*.

Step 1: NCVS's contextual information acquisition module automatically collects comprehensive dependency information, including topology, service deployment, and correlations between components. Then, the module stores the collected information into the contextual information DB for post processing. All the operations within this step do not require any additional instrumentations from administrators. §3.2 discusses this module's design in more detail.

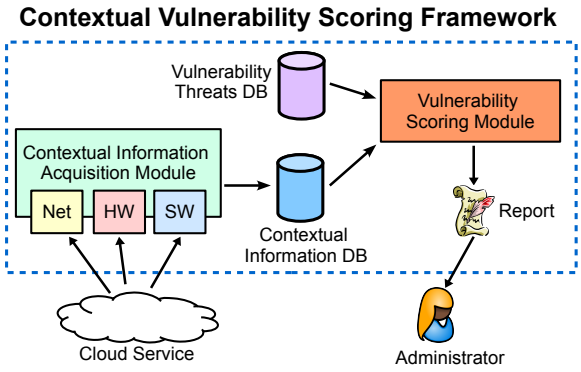

Figure 2: NCVS's architecture overview.

Step 2: The vulnerability scoring module takes the information in the contextual information DB and vulnerability threats DB as input, and constructs a dependency graph, named CDG, modeling the context of the target service *S*. While various graphs representing service structures have been proposed in the past years, CDG is the first effort capable of capturing comprehensive contextual information. With the CDG in hand, the module analyzes it and outputs a report ranking vulnerabilities based on their scores specific to the context of the target service *S*. §3.3 details this module.

3.2 Contextual Information Acquisition

For a target service *S*, NCVS's contextual information acquisition module automatically collects three types of dependency information (network dependency, hardware dependency and software dependency), and then stores the collected information into the contextual information database (as shown in Figure 2). We focus on the collection of these three types of contextual information, since this information is able to comprehensively represent a service's context (or environment) [30, 31]. To our knowledge, there is no existing effort that can simultaneously collect these three types of dependency information without any instrumentation.

Hardware dependency means *S*'s physical topology information, including the locations of servers and network devices (*e.g.*, aggregation switches and core routers) as well as the links between servers and network devices. In other words, hardware dependency information is the data-center infrastructure components and their links used to support service *S* [35].

Software dependency denotes software components running on each node (*e.g.*, server or switch). For switches, the software components might be the switch OS and related embed functional applications; for servers, the software components include operating systems, service *S*'s components (*e.g.*, back-end DB), and needed applications (*e.g.*, SSH) [17, 25–28, 32, 36].

Network dependency means an invoking sequence (or flow)

between different software components in the target service *S*. For example in a MapReduce service, a request is first sent to a job tracker and then arrives at several task trackers, forming a directed invoking flow. Such an invoking sequence, *e.g.*, Client -> JobTracker -> TaskTrackers, represents an item of network dependency information. Different from hardware dependency, network dependency is at logical level, which means a network dependency (*i.e.*, an invoking flow) may be across different servers or only occurs on the same server.

As described in the introduction section, the first challenge of building NCVS is how to acquire all of the above three types of information without any instruments or modifications to the source code of target services. We now describe how we develop the contextual information acquisition module.

Software dependency collection. The basic idea of software dependency collector is to determine software components running on each node (*e.g.*, server or switch) by analyzing the node's log information during a service request session. A software dependency is defined as `<node="M" sw="S" dep="x,y,z"/>`, where `node` indicates the node running (or holding) this software component specified by `sw`, and `dep` shows all the components (*e.g.*, libraries) used by this software component `sw`.

Our insight for designing the software dependency collector is: when a request is sent to the target service *S*, each involved node would generate log corresponding to its functions, and the log information across all the nodes has a time-aware sequence. Thus, if we can capture and analyze all the involved nodes' log information between the start time and end time of accessing *S*, we would be able to extract the software dependency information of each node supporting *S*. In our design, therefore, the software dependency collector – playing as a client role – sends many requests to the target service *S* and infers software dependencies of each node by analyzing the generated log during the session. The collector employs a well-known association rule mining algorithm, named Apriori [6], that first groups each item of log information by a time interval, then mines correlated rules among log items with confidence values, and finally selects these correlations with the high confidence level as software dependencies. The accuracy of collector is related to the threshold of time interval we set – smaller threshold gives better accuracy.

Network dependency collection. In order to automatically acquire needed network dependency information, we design a network dependency collector based on the similar intuition as the software dependency collector. In particular, we define a network dependency as a continuous, directed stream of packets between two services (or applications). Similar to the software dependency collector, the network dependency collector also sends many requests to the target service *S*, but it aims to capture all the networking packets and flows during the service access session. To capture these traffics, our collector employs NSDMiner [14, 20], which is a well-developed traffic monitor. An important reason we choose NSDMiner is that NSDMiner does not need to install extra agents or software on involved nodes; moreover, NSDMiner can collect more accurate traffic, thus making our network dependency collector work better than Sherlock [8] and Orion [10], two representative network dependency collection tools. The network dependency collector outputs many items of the network dependency information formatted as follows:

```
<IP address A>:<port A>:[TCP] <# App~1>
  <IP address B>  <port B>  [TCP]
  <IP address D>  <port D>  [TCP]
<IP address B>:<port B>:[UDP] <# App~2>
  <IP address C>  <port C>  [UDP]
  <IP address D>  <port D>  [UDP]
...
```

This example shows that component *A* which runs on IP address *A* at port *A* depends on components *B* and *D* which run on their respective IP addresses and ports. Similarly, component *B* depends on component *C* and *D*. Note that all the components (A, B, C and D) are services or applications.

Hardware dependency collection. Since most of current data-center networks employ software defined network (SDN) controller [18] or network provenance techniques [39, 40], it is easy for the administrators to learn physical network topology (*i.e.*, hardware dependency information for our system). Thus, if an administrator has known network topology, our hardware dependency collector would directly read and parse the network topology file that might be written in different formats (*e.g.*, CSV or XML), thus obtaining our needed hardware dependency information.

In the case that there is no such file available, we also provide a network topology generation toolkit (TopoGen) to assist administrators to derive the network topology in a fast, less tedious way. TopoGen supports various data-center topology models including fat tree [19] and BCube [12]. The administrators only need to give the basic information, *i.e.*, the number of core/aggregation/edge switches, to TopoGen, and then TopoGen generates a network topology automatically. Administrators are allowed to validate the auto-generated topology and further modify if necessary. Furthermore, TopoGen allows users to plug new topology generation rules into generate network links, which is very efficient and helpful for large-scale network systems. Thus, combining with our TopoGen toolkit and the basic information from cloud administrators, we can derive accurate network topology information in a less tedious manner.

3.3 Vulnerability Scoring

Vulnerability scoring module computes the contextual score of each vulnerability in *S* by two steps: 1) modeling *S*'s context with the contextual information acquired in §3.2, and 2) computing the importance of each vulnerability in terms of this model.

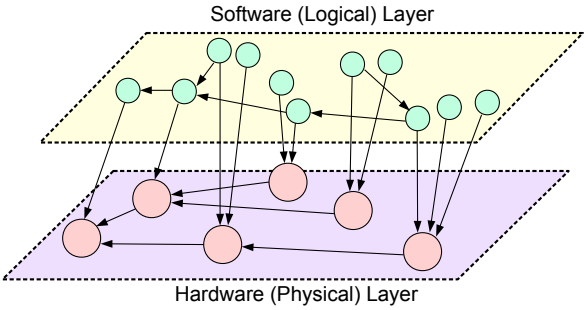

Figure 3: Contextual dependency graph (CDG).

3.3.1 Contextual Dependency Graph (CDG)

The vulnerability scoring module first needs to build an explicit graph, named contextual dependency graph (or CDG), to represent the context of the target service S. Figure 3 presents a CDG example.

In a CDG, there are two layers: hardware (physical) layer and software (logical) layer. All the nodes representing physical components (*e.g.*, servers and switches) are involved in the hardware layer; on the contrary, all the nodes representing software components (*e.g.*, libraries and applications) are put in the software layer. All the edges in the CDG are directed. An edge like $A \rightarrow B$ means component A depends on another component B.

Because our contextual information acquisition module automatically collects hardware, software and network dependency information, we now describe the correlation between the collected contextual information and a CDG. All the hardware components (*e.g.*, servers and switches) and their dependencies (*e.g.*, links) are modeled as CDG nodes in the hardware layer and directed edges among them, respectively. In other words, a directed edge between two CDG nodes in the hardware layer means a physical topology link. All the software dependencies of each machine are modeled as CDG nodes in the software layer, and they are connected by directed edges to their corresponding nodes in the hardware layer, as shown in Figure 3. Namely, a directed edge across the hardware layer and the software layer in a CDG means a software dependency. All the network dependencies are modeled as directed edges connecting different CDG nodes in the software layer. Namely, a directed edge between two CDG nodes in the software layer means a network dependency.

3.3.2 Vulnerabilities' Scores Computation

With a CDG in hand, the vulnerability scoring module now aims to identify what nodes have vulnerabilities in the CDG and score these identified vulnerabilities by taking into account their "roles" in this CDG. As shown in Figure 2, the vulnerability scoring module relies on another database, named vulnerability threats DB. The vulnerability threats DB could be any existing vulnerability scoring system – our prototype employs CVSS. We need such a database because it can offer exploited vulnerabilities and their basic threats given by security experts.

Identify involved vulnerabilities. The vulnerability scoring module first uses the node information in CDG to search vulnerabilities involved in the service S's deployment. In our prototype, we implemented keyword matching and classification methods to identify involved vulnerabilities in CVSS, since CVSS provides detailed information about each recorded vulnerability such as its software name, threat, and impact. Note that because almost of all the vulnerability databases store software vulnerabilities, NCVS also focuses on software vulnerabilities rather than hardware vulnerabilities.

Node ranking by considering context. After determining the involved vulnerabilities in the CDG, we start to score nodes and vulnerabilities in our target context. Specifically, for each affected software node n in CDG, we compute its importance in terms of three different sub-graphs in CDG:

- Topology-aware importance $ti(n) = Rank(hw_graph, n)$, where hw_graph means the nodes that are located at the hardware layer and have software component n.
- Software-aware importance $si(n) = Rank(sw_graph, n)$, where sw_graph means the nodes that are located at the software layer and have software component n.
- Network-aware important $ni(n) = Rank(net_graph, n)$, where net_graph means the nodes that are involved in network dependencies and have software component n.

In the above computation, $Rank()$ means graph node ranking function. In our prototype, we adopt PageRank algorithm [23] to enable the graph node ranking function, *i.e.*, computing the importance of the node n. Our prototype supports any graph node ranking algorithm, including PageRank, HITs [16] and new developed ranking algorithms.

Ranking vulnerabilities. With the importance of each software node in hand, we can compute the contextual scores for vulnerabilities involved in our target service S. In particular, for a given vulnerability v, we compute the score of v, $severity(v)$, as follows.

$$severity(v) = w_{ti} * \sum_{i=1}^{|S_v|} ti(n_i) + w_{ni} * \sum_{i=1}^{|S_v|} ni(n_i) + w_{si} * \sum_{i=1}^{|S_v|} si(n_i)$$

Where, S_v means the set of software nodes that contain the vulnerability v. w_{ti}, w_{ni} and w_{si} are used to weight the three types of context impacts, respectively. These weights are first assigned based on expert knowledge and further tuned in an iterative process. For example, we can start with equal importance for three contexts $w_{ti} = w_{ni} = wli = 1$ and adjust these weights to control the score until the final scores are reasonable. The design of the aggregation function is also pluggable, which means developers can customize their aggregate function according to their particular purpose. For example, in our prototype, we also provide another aggregate function to integrate CVSS score based on the product rule: $severity(v) = (\sum_{i=1}^{|S_v|} ti(n_i) * ni(n_i) * si(n_i)) * CVSS(v)$. In this aggregate function, CVSS score is considered as a local software importance factor.

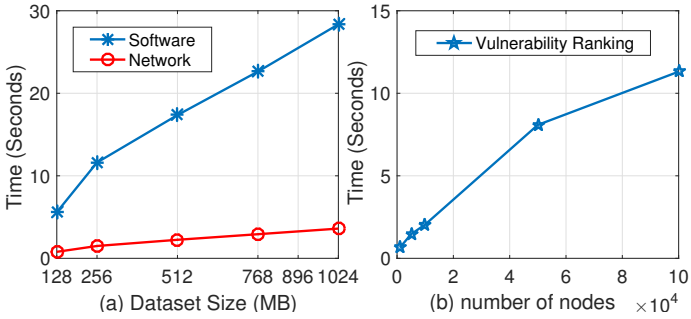

Figure 4: Performance evaluation: (a) contextual information acquisition module, and (b) vulnerability scoring module.

Table 1: Comparing the vulnerability scores output by CVSS and NCVS in our case study. The fourth column in this table contains scores produced by CVSS, while the fifth column includes NCVS's scores for our evaluated vulnerabilities.

Seq	ID	Type	CVSS	NCVS
1	CVE-2015-1776	Hadoop	6.3	0.0257
2	CVE-2015-4279	Network	7.8	0.1069
3	CVE-2015-5210	Hadoop	5.8	0.0275
4	CVE-2015-6355	Network	5.0	0.0403
5	CVE-2015-6415	Network	7.1	0.0770
6	CVE-2015-6420	Hadoop	7.5	0.0287
7	CVE-2015-7430	Hadoop	8.4	0.0293
8	CVE-2016-0731	Hadoop	4.9	0.0269
9	CVE-2016-1503	Network	9.8	0.0789
10	CVE-2016-2170	Hadoop	9.8	0.0303

4. EVALUATION

We have developed a prototype system in Java to evaluate NCVS. In this section, we first evaluate NCVS's performance (§4.1), and then evaluate the effectiveness of NCVS with a real-world case study (§4.2).

4.1 Performance Evaluation

We deploy our NCVS prototype on a workstation equipped with Intel Xeon E5-2630 2.30 GHz CPU and 64 GB RAM. For performance evaluation, we measure the performance of two modules of our NCVS prototype – *i.e.*, the runtime for two modules to handle different sizes of services.

First, we evaluate the runtime of the first module of NCVS, *i.e.*, contextual information acquisition module. The performance bottlenecks of this module are software dependency collector and network dependency collector, so that our experiments focus on these two collectors.

Figure 4(a) presents the evaluational results for software and network dependency collectors. In particular, we vary the size of target services, thus getting raw datasets with different sizes (reflected by the x-axis of Figure 4(a)). Then, we measured how long two modules need to extract dependency information of interest. As shown in Figure 4(a), the runtime of two collectors increases with the dataset size, but the software dependency collector is more expensive due to the complexity of Apriori algorithm. It is worth mentioning that the increase in runtime over the size of services is acceptable in practice, since such a process is typically operated off-line and the results could be re-used in the future.

On the right hand, Figure 4(b) shows that the runtime of vulnerability ranking algorithm (the performance bottleneck of the vulnerability scoring module) increases with the number of nodes in the CDG. We can observe that NCVS is quite efficient, with less than 12 seconds ranking 100,000 nodes. Our ranking algorithm is powered by PageRank algorithm which is highly efficient for large-scale graphs (here we set the maximum number of iteration is 100 and the error tolerance of two consecutive iterations is 0.001).

4.2 Case Study

In order to evaluate the effectiveness of NCVS, we construct a real-world case study similar to our motivating example (in §2.1). In the case study, we deploy a Hadoop service on a lab-scale cluster with 16 server machines and 8 switches. The CDG generated by NCVS to model this service contains 747 nodes and 1204 dependencies. As shown in Table 1, we extract ten vulnerabilities from this service and derive their scores from CVSS (the fourth column in Table 1). We observe that the vulnerabilities 2, 4, 5, 9 are network-related and located in the switches, while the remaining ones are vulnerabilities of Hadoop and its related libraries (*e.g.*, Apache common library). The fifth column in Table 1 shows these ten vulnerabilities' scores produced by NCVS.

In Table 1, we observe that if we fix these vulnerabilities based on the scores output by CVSS, the order should be 10 -> 9 -> 7 -> 2 -> 6 -> 5 -> 1 -> 3 -> 4 -> 8. On the contrary, if we fix vulnerabilities according to the scores output by NCVS, the order should be 2 -> 9 -> 5 -> 4 -> 10 -> 7 -> 6 -> 3 -> 8 -> 1. To demonstrate NCVS's effectiveness, we construct an interesting "vulnerability fixing" experiment. We assume that the current cluster is compromised by these ten vulnerabilities and all 16 server machines are not alive. Now, we fix these vulnerabilities based on the above two orders output by CVSS and NCVS, respectively. At each time point t, we only fix one vulnerability. After each time point, we evaluate the system in terms of the number of alive nodes in the cluster (as a safety metric). Figure 5 shows our experimental results. We observe that fixing vulnerabilities according to the scores output by NCVS can recover the capability of the Hadoop cluster much faster than fixing vulnerabilities based on CVSS ranking list. Thus, we can say NCVS can output more meaningful vulnerabilities' scores than CVSS since NCVS takes into account the context of the target service.

In addition, NCVS's results have the following advantages. First, the result is network topology aware. We observe that network-related vulnerabilities in NCVS's ranking are generally higher than that of CVSS ranking list. Second, the result can distinguish the roles of different software components. We observe that the vulnerability (*i.e.*, CVE-

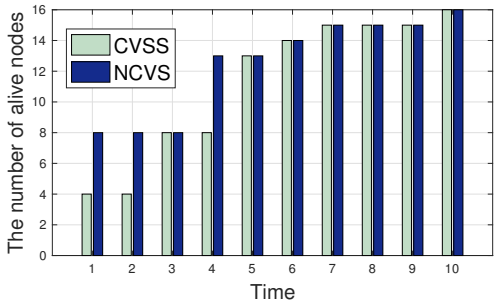

Figure 5: Comparison of two different fix plans.

2015-7430) that affect master nodes ranks higher than that that affect slave nodes (*i.e.*, CVE-2015-1776). Finally, the result is accurate. In the case that two vulnerabilities are located in the same software component (*e.g.*, CVE-2015-6420 and CVE-2016-2170), their severity ordering depend on their scores in CVSS, which means NCVS does not violate the basic severity of each vulnerability.

5. CONCLUSION

In this paper, we propose, design, implement and evaluate a novel framework, NCVS, that can automatically score vulnerabilities for cloud services. Different from existing efforts, NCVS's workflow is non-intrusive and NCVS scores any given vulnerability by not only considering its intrinsic threats, but also taking into account its service context. Our evaluational results demonstrate NCVS's effectiveness (by comparing with CVSS) and performance.

6. REFERENCES

[1] Common Vulnerabilities and Exposures (CVE). `https://cve.mitre.org/`.

[2] Common Vulnerability Scoring System (CVSS). `http://nvd.nist.gov/cvss.cfm`.

[3] Common Vulnerability Scoring System v3.0 (CVSS 3.0). `https://www.first.org/cvss/specification-document`.

[4] Open Vulnerability and Assessment (OVAL). `http://oval.mitre.org/`.

[5] OSVDB. The Open Source Vulnerability Database. `http://osvdb.org/`.

[6] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining association rules between sets of items in large databases. *ACM SIGMOD Record*, 22(2):207–216, 1993.

[7] Marcos Kawazoe Aguilera, Jeffrey C. Mogul, Janet L. Wiener, Patrick Reynolds, and Athicha Muthitacharoen. Performance debugging for distributed systems of black boxes. In *19th SOSP*, October 2003.

[8] Paramvir Bahl, Ranveer Chandra, Albert G. Greenberg, Srikanth Kandula, David A. Maltz, and Ming Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *SIGCOMM*, August 2007.

[9] Paul Barham, Austin Donnelly, Rebecca Isaacs, and Richard Mortier. Using Magpie for request extraction and workload modelling. In *6th OSDI*, December 2004.

[10] Xu Chen, Ming Zhang, Zhuoqing Morley Mao, and Paramvir Bahl. Automating network application dependency discovery: Experiences, limitations, and new solutions. In *8th OSDI*, December 2008.

[11] Christian Frühwirth and Tomi Männistö. Improving CVSS-based vulnerability prioritization and response with context information. In *3rd ESEM*, October 2009.

[12] Chuanxiong Guo, Guohan Lu, Dan Li, Haitao Wu, Xuan Zhang, Yunfeng Shi, Chen Tian, Yongguang Zhang, and Songwu Lu. Bcube: a high performance, server-centric network architecture for modular data centers. In *SIGCOMM*, August 2009.

[13] Heqing Huang, Su Zhang, Xinming Ou, Atul Prakash, and Karem A. Sakallah. Distilling critical attack graph surface iteratively through minimum-cost SAT solving. In *27th ACSAC*, December 2011.

[14] Barry Peddycord III, Peng Ning, and Sushil Jajodia. On the accurate identification of network service dependencies in distributed systems. In *26th LISA*, December 2012.

[15] Jianchun Jiang, Liping Ding, Ennan Zhai, and Ting Yu. VRank: A context-aware approach to vulnerability scoring and ranking in SOA. In *6th SERE*, June 2012.

[16] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. In *9th SODA*, January 1998.

[17] Bo Liu, Ennan Zhai, Huiping Sun, Yelu Chen, and Zhong Chen. Filtering spam in social tagging system with dynamic behavior analysis. In *2009 International Conference on Advances in Social Network Analysis and Mining, ASONAM 2009, 20-22 July 2009, Athens, Greece*, pages 95–100, 2009.

[18] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. OpenFlow: Enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, 2008.

[19] Radhika Niranjan Mysore, Andreas Pamboris, Nathan Farrington, Nelson Huang, Pardis Miri, Sivasankar Radhakrishnan, Vikram Subramanya, and Amin Vahdat. PortLand: A scalable fault-tolerant layer 2 data center network fabric. In *SIGCOMM*, August 2009.

[20] Arun Natarajan, Peng Ning, Yao Liu, Sushil Jajodia, and Steve E. Hutchinson. NSDMiner: Automated discovery of network service dependencies. In *31st INFOCOM*, March 2012.

[21] Xinming Ou, Wayne F. Boyer, and Miles A. McQueen. A scalable approach to attack graph generation. In *13th CCS*, November 2006.

[22] Xinming Ou, Sudhakar Govindavajhala, and Andrew W. Appel. MulVAL: A logic-based network security analyzer. In *14th USENIX Security*, July 2005.

[23] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank citation ranking: Bringing order to the web. 1999.

[24] Reginald E. Sawilla and Xinming Ou. Identifying critical attack assets in dependency attack graphs. In *13th ESORICS*, October 2008.

[25] Cong Sun, Ennan Zhai, Zhong Chen, and Jianfeng Ma. A multi-compositional enforcement on information flow security. In *Information and Communications Security - 13th International Conference, ICICS 2011, Beijing, China, November 23-26, 2011. Proceedings*, pages 345–359, 2011.

[26] Yonggang Wang, Ennan Zhai, Cui Cao, Yongqiang Xie, Zhaojun Wang, Jian-bin Hu, and Zhong Chen. Dspam: Defending against spam in tagging systems via users' reliability. In *16th IEEE International Conference on Parallel and Distributed Systems, ICPADS 2010, Shanghai, China, December 8-10, 2010*, pages 139–146, 2010.

[27] Yonggang Wang, Ennan Zhai, Jian-bin Hu, and Zhong Chen. Claper: Recommend classical papers to beginners. In *Seventh International Conference on Fuzzy Systems and Knowledge Discovery, FSKD 2010, 10-12 August 2010, Yantai, Shandong, China*, pages 2777–2781, 2010.

[28] Yonggang Wang, Ennan Zhai, Eng Keong Lua, Jian-bin Hu, and Zhong Chen. isac: Intimacy based access control for social network sites. In *9th International Conference on Ubiquitous Intelligence and Computing and 9th International Conference on Autonomic and Trusted Computing, UIC/ATC 2012, Fukuoka, Japan, September 4-7, 2012*, pages 517–524, 2012.

[29] Ennan Zhai, Ruichuan Chen, Zhuhua Cai, Long Zhang, Eng Keong Lua, Huiping Sun, Sihan Qing, Liyong Tang, and Zhong Chen. Sorcery: Could we make P2P content sharing systems robust to deceivers? In *Proceedings P2P 2009, Ninth International Conference on Peer-to-Peer Computing, 9-11 September 2009, Seattle, Washington, USA*, pages 11–20, 2009.

[30] Ennan Zhai, Ruichuan Chen, David Isaac Wolinsky, and Bryan Ford. An untold story of redundant clouds: making your service deployment truly reliable. In *Proceedings of the 9th Workshop on Hot Topics in Dependable Systems, HotDep 2013, Farmington, Pennsylvania, USA, November 3, 2013*, pages 3:1–3:6, 2013.

[31] Ennan Zhai, Ruichuan Chen, David Isaac Wolinsky, and Bryan Ford. Heading off correlated failures through Independence-as-a-service. In *11th OSDI*, October 2014.

[32] Ennan Zhai, Liping Ding, and Sihan Qing. Towards a reliable spam-proof tagging system. In *Fifth International Conference on Secure Software Integration and Reliability Improvement, SSIRI 2011, 27-29 June, 2011, Jeju Island, Korea*, pages 174–181, 2011.

[33] Ennan Zhai, Liang Gu, and Yumei Hai. A risk-evaluation assisted system for service selection. In *2015 IEEE International Conference on Web Services, ICWS 2015, New York, NY, USA, June 27 - July 2, 2015*, pages 671–678, 2015.

[34] Ennan Zhai, Zhenhua Li, Zhenyu Li, Fan Wu, and Guihai Chen. Resisting tag spam by leveraging implicit user behaviors. *PVLDB*, 10(3):241–252, 2016.

[35] Ennan Zhai, Qingni Shen, Yonggang Wang, Tao Yang, Liping Ding, and Sihan Qing. Secguard: Secure and practical integrity protection model for operating systems. In *Web Technologies and Applications - 13th Asia-Pacific Web Conference, APWeb 2011, Beijing, China, April 18-20, 2011. Proceedings*, pages 370–375, 2011.

[36] Ennan Zhai, Huiping Sun, Sihan Qing, and Zhong Chen. Spamclean: Towards spam-free tagging systems. In *Proceedings of the 12th IEEE International Conference on Computational Science and Engineering, CSE 2009, Vancouver, BC, Canada, August 29-31, 2009*, pages 429–435, 2009.

[37] Ennan Zhai, Huiping Sun, Sihan Qing, and Zhong Chen. Sorcery: Overcoming deceptive votes in P2P content sharing systems. *Peer-to-Peer Networking and Applications*, 4(2):178–191, 2011.

[38] Xu Zhao, Yongle Zhang, David Lion, Muhammad Faizan Ullah, Yu Luo, Ding Yuan, and Michael Stumm. lprof: A non-intrusive request flow profiler for distributed systems. In *11th OSDI*, October 2014.

[39] Wenchao Zhou, Qiong Fei, Arjun Narayan, Andreas Haeberlen, Boon Thau Loo, and Micah Sherr. Secure network provenance. In *23rd SOSP*, October 2011.

[40] Wenchao Zhou, Qiong Fei, Shengzhi Sun, Tao Tao, Andreas Haeberlen, Zachary G. Ives, Boon Thau Loo, and Micah Sherr. NetTrails: a declarative platform for maintaining and querying provenance in distributed systems. In *SIGMOD*, June 2011.